arXiv:1704.08676v1 [cs.LG] 27 Apr 2017

A quantitative assessment of the effect of different algorithmic schemes to the task of learning the structure of Bayesian Networks

Stefano Beretta[a], Mauro Castelli[b], Ivo Gonçalves[b,c], Daniele Ramazzotti[d,*]

[a] *DISCo, Universitá degli Studi di Milano-Bicocca. 20126 Milano, Italy*
[b] *NOVA IMS, Universidade Nova de Lisboa. 1070-312 Lisboa, Portugal*
[c] *CISUC, Department of Informatics Engineering, University of Coimbra, 3030-290 Coimbra, Portugal*
[d] *Department of Pathology, Stanford University. California, USA*

Abstract

One of the most challenging tasks when adopting Bayesian Networks (BNs) is the one of learning their structure from data. This task is complicated by the huge search space of possible solutions and turned out to be a well-known NP-hard problem and, hence, approximations are required. However, to the best of our knowledge, a quantitative analysis of the performance and characteristics of the different heuristics to solve this problem has never been done before.

For this reason, in this work, we provide a detailed study of the different state-of-the-arts methods for structural learning on simulated data considering both BNs with discrete and continuous variables, and with different rates of noise in the data. In particular, we investigate the characteristics of different widespread scores proposed for the inference and the statistical pitfalls within them.

Keywords: Bayesian Networks, Structure Learning, Heuristic Search, Evolutionary Computation, Genetic Algorithms

*Corresponding author
Email addresses: stefano.beretta@disco.unimib.it (Stefano Beretta), mcastelli@novaims.unl.pt (Mauro Castelli), igoncalves@novaims.unl.pt (Ivo Gonçalves), daniele.ramazzotti@stanford.edu (Daniele Ramazzotti)

1. Introduction

Bayesian Networks (BNs) have been applied to several different fields, ranging from the water resources management [1] to the discovery of gene regulatory networks [2, 3]. The task of learning a BN can be divided into two subtasks: (1) structural learning, i.e., identification of the topology of the BN, and (2) parametric learning, i.e., estimation of the numerical parameters (conditional probabilities) for a given network topology. In particular, the most challenging task of the two is the one of learning the structure of a BN. Different methods have been proposed to face this problem, and they can be classified into two categories [4, 5]: (1) methods based on detecting conditional independencies, also known as constraint-based methods, and (2) score+search methods, also known as score-based approaches.

As discussed in [6], the input of the former algorithms is a set of conditional independence relations between subsets of variables, which are used to build a BN that represents a large percentage (and, whenever possible, all) of these relations [7]. However, the number of conditional independence tests that such methods should perform is exponential and, thus, approximation techniques are required.

Although constraint-based learning is an interesting approach, as it is close to the semantic of BNs, most of the developed structure learning algorithms fall into the score-based method category, given the possibility of formulating such a task in terms of an optimization problem. As the name implies, these methods have two major components: (1) a scoring metric that measures the quality of every candidate BN with respect to a dataset, and (2) a search procedure to intelligently move through the space of possible networks, as this space is enormous. More in detail, as shown in [8, 9], searching this space for the optimal structure is an NP-hard problem, even when the maximum number of parents for each node is constrained.

Hence, regardless of the strategy to learn the structure, one wishes to pursue, greedy local search techniques and heuristic search approaches need to be adopted to tackle the problem of inferring the structure of BNs. However, to the best of our knowledge, a quantitative analysis of the performance of these different techniques has never been done. For this reason, here we provide a detailed assessment of the performance of different state-of-the-arts methods for structural learning on simulated data, considering both BNs with discrete and continuous variables, and with different rates of noise in the data. More precisely, we investigate the characteristics of different

scores proposed for the inference, and the statistical pitfalls within both the constraint-based and score based techniques. Furthermore, we study the impact of various heuristic search approaches, namely Hill Climbing, Tabu Search, and Genetic Algorithms.

This work is structured as follows. In the next two Sections, we provide a background on both Bayesian Networks and Heuristic Search Techniques (see Section 2 and Section 3). In Section 4 we describe the results of our study and in Section 5 we conclude the paper.

2. Bayesian Networks

BNs are graphical models representing the *joint distribution* over a set of n random variables through a *direct acyclic graph* (DAG) $G = (V, E)$, where the n nodes in V represent the variables, and the arcs E encode any statistical dependence among them. In this DAG, the set of variables with an arc toward a given node X is its set $\pi(X)$ of "parents". When the structure of a BN is known, it is possible to compute the joint distribution of all the variables as the product of the conditional distributions on each variable given its parents. However, even if we consider the simplest case of binary variables, the number of parameters in each conditional probability table is still exponential in size. For example, in the case of binary variables, the total number of parameters needed to compute the full joint distribution is of size $\sum_X 2^{|\pi(X)|} - 1$, where $|\pi(X)|$ is the cardinality of the parent set of each node.

In the literature, there are two families of methods to learn the structure of a BN from data. The idea behind the first group of methods is the one of *learning the conditional independence relations* of the BN from which, in turn, the network is learned. These methods are often referred to as *constraint-based approaches*. The second group of methods, the so-called *score-based approaches*, formulates the task of structure learning as an optimization problem, with scores aimed at *maximizing the likelihood of the data* given the model. However, both the approaches are known to lead to NP-hard formulations and, because of this, heuristic methods need to be used to find a "good" solution.

2.1. Constraint-based approaches

We now briefly describe the main idea behind this class of approaches. For a detailed discussion of this topic we refer to [7, 10].

This class of methods aims at building a graph structure to reflect the dependence relations in the data that match the empirical distribution. Notwithstanding, the number of conditional independence tests that such algorithms would have to perform among any pair of nodes to test all possible relations is exponential and, because of this, the introduction of some approximations is required.

We now provide some details on two constraint-based algorithms that have been proven to be efficient and of widespread usage: the *PC algorithm* [7] and the *Incremental Association Markov Blanket* (IAMB) [10].

The PC algorithm. This algorithm [7] starts with a fully connected graph and, on the basis of pairwise independence tests, it iteratively removes all the extraneous edges. To avoid an exhaustive search of separating sets, the edges are ordered to consider the correct ones early in the search. Once a separating set is found, the search for that pair ends. The PC algorithm orders the separating sets by increasing values of size l, starting from 0 (the empty set), until $l = n - 2$ (where n is the number of variables). The algorithm stops when every variable has less than $l - 1$ neighbors since it can be proven that all valid sets must have already been chosen. During the computation, the bigger the value of l is, the larger the number of separating sets must be considered. However, by the time l gets too big, the number of nodes with degree l or higher must have dwindled considerably. Thus, in practice, we only need to consider a small subset of all the possible separating sets.

Incremental Association Markov Blanket algorithm. A different type of constraint-based learning algorithms uses the Markov blankets [11] to restrict the subset of variables to test for independence. Thus, when this knowledge is available in advance, we do not need to test a conditioning on all possible variables. A widely used and efficient algorithm for Markov blanket discovery is IAMB which, for each variable X, keeps track of a hypothesis set $\mathcal{H}(X)$. The goal is, for a given $\mathcal{H}(X)$, to obtain at the end of the algorithm, a Markov blanket of X equal to $\mathcal{B}(X)$. IAMB consists of a forward and a backward phase. During the forward phase, it adds all the possible variables into $\mathcal{H}(X)$ that could be in $\mathcal{B}(X)$, while in the backward phase, it removes all the false positive variables from the hypotheses set, leaving the true $\mathcal{B}(X)$. The forward phase begins with an empty $\mathcal{H}(X)$ for each X. Then, iteratively, variables with a strong association with X (conditioned on all the variables in $\mathcal{H}(X)$) are

added to the hypotheses set. This association can be measured by a variety of non-negative functions, such as *mutual information*. As $\mathcal{H}(X)$ grows large enough to include $\mathcal{B}(X)$, the other variables in the network will have very little association with X, conditioned on $\mathcal{H}(X)$. At this point, the forward phase is complete. The backward phase starts with $\mathcal{H}(X)$ that contains $\mathcal{B}(X)$ and false positives, which will have small conditional associations, while true positives will associate strongly. Using this test, the backward phase is able to iteratively remove the false positives, until all but the true negatives are eliminated.

2.2. Score-based approaches

These approaches aim at maximizing the likelihood $\mathcal{L}$ of a set of observed data D, which can be computed as the product of the probability of each observation. Formally,

$$\mathcal{L}(D) = \prod_{d \in D} P(d)$$

for a set of observations D. Since we want to infer a model G that best explains the observed data, we define the likelihood of observing the data given a specific model G as:

$$\mathcal{LL}(G, D) = \prod_{d \in D} P(d|G)\,.$$

Practically, however, for any arbitrary set of data, the most likely graph is always the fully connected one, since adding an edge can only increase the likelihood of the data, i.e., this approach overfits the data. To overcome this limitation, the likelihood score is almost always combined with a *regularization term* that penalizes the complexity of the model in favor of sparser solutions [11].

As already mentioned, such an optimization problem leads to intractability, due to the enormous search space of the valid solutions. Because of this, the optimization task is often solved with heuristic techniques. Before moving on to describe the main heuristic methods employed to face such complexity (see Section 3), we now give a short description of a particularly relevant and known score, called *Bayesian Information Criterion* (BIC) [12], as an example of scoring function adopted by several the score-based methods.

The Bayesian Information Criterion. BIC uses a score that consists of a log-likelihood term and a regularization term that depend on a model G and data D:

$$\text{BIC}(G, D) = \mathcal{LL}(G, D) - \frac{\log m}{2}\dim(G) \tag{1}$$

where D denotes the data, m denotes the number of samples, and $\dim(G)$ denotes the number of parameters in the model. Since, in general, $\dim(\cdot)$ depends on the number of parents of each node, it is a good metric for model complexity. Moreover, each edge added to G increases the complexity of the model. Thus, the regularization term based on $\dim(\cdot)$ favors graphs with fewer edges and, more specifically, fewer parents for each node. The term $\log m/2$ essentially weighs the regularization term. The effect is that the higher the weight, the more sparsity will be favored over "explaining" the data through maximum likelihood.

Notice that the likelihood is implicitly weighted by the number of data points since each point contributes to the score. As the sample size increases, both the weight of the regularization term and the "weight" of the likelihood increase. However, the weight of the likelihood increases faster than that of the regularization term. This means that, with more data, the likelihood will contribute more to the score, and we may trust our observations more and have less need for regularization. Statistically speaking, BIC is a *consistent score* [11]. Consequently, G contains the same independence relations as those implied by the true structure.

3. Heuristic Search Techniques

We now describe some of the main state-of-the-art search strategies that we took into account in this work. In particular, as stated in Section 1, we considered the following search methods: Hill Climbing, Tabu Search, and Genetic Algorithms.

3.1. Hill Climbing

Hill Climbing (HC) is one of the simplest iterative techniques that have been proposed for solving optimization problems. While HC consists of a simple and intuitive sequence of steps, it is a good search scheme to be used as a baseline for comparing the performance of more advanced optimization techniques.

Hill Climbing shares with other techniques (like Simulated Annealing [13] and Tabu Search [14]) the concept of neighborhood. Search methods based on this latter concept are iterative procedures in which a neighborhood $N(i)$ is defined for each feasible solution i, and the next solution j is searched among the solutions in $N(i)$. Hence, the neighborhood is a function $N : S \rightarrow 2^S$ that assigns to each solution in the search space S a (non-empty) subset of S.

The sequence of steps of the Hill Climbing algorithm, for a minimization problem w.r.t. a given objective function f, are the following:

1. choose an initial solution i in S;
2. find the best solution j in $N(i)$ (i.e., the solution j such that $f(j) \leq f(k)$ for every k in $N(i)$;
3. if $f(j) > f(i)$, then stop; else set $i = j$ and go to Step 2.

As it is clear from the aforementioned algorithm, Hill Climbing returns a solution that is a local minimum for the problem at hand. This local minimum does not generally correspond to the global minimum for the problem under exam, that is, Hill Climbing does not guarantee to return the best possible solution for a given problem. To counteract this limitation, more advanced neighborhood search methods have been defined. One of these methods is Tabu Search, an optimization technique that uses the concept of "memory".

3.2. Tabu Search

Tabu Search (TS) is a meta-heuristic that guides a local heuristic search procedure to explore the solution space beyond local optimality. One of the main components of this method is the use of an adaptive memory, which creates a more flexible search behavior. Memory-based strategies are therefore the main feature of TS approaches, founded on a quest for "integrating principles", by which alternative forms of memory are appropriately combined with effective strategies for exploiting them.

Tabus are one of the distinctive elements of TS when compared to Hill Climbing or other local search methods. The main idea in considering *tabus* is to prevent cycling when moving away from local optima through non-improving moves. When this situation occurs, something needs to be done to prevent the search from tracing back its steps to where it came from. This is achieved by declaring *tabu* (disallowing) moves that reverse the effect of recent moves. For instance, let us consider a problem where solutions are

binary strings of a prefixed length, and the neighborhood of a solution i consists of the solutions that can be obtained from i by flipping only one of its bits. In this scenario, if a solution j has been obtained from a solution i by changing one bit b, it is possible to declare a *tabu* to avoid to flip back the same bit b of j for a certain number of iterations (this number is called the tabu tenure of the move). *Tabus* are also useful to help in moving the search away from previously visited portions of the search space and, thus, perform more extensive exploration.

As reported in [15], tabus are stored in a short-term memory of the search (the tabu list) and usually, only a fixed limited quantity of information is recorded. It is possible to store complete solutions, but this has a negative impact on the computational time required to check whether a move is a tabu or not and, moreover, it requires a vast amount of space. The second option (which is the one commonly used) involves recording the last few transformations performed to obtain the current solution and prohibiting reverse transformations.

While tabus represent the main distinguish feature of TS, this feature can introduce other issues in the search process. In particular, the use of tabus can prohibit attractive moves, or it may lead to an overall stagnation of the search process, with a lot of moves that are not allowed. Hence, several methods for revoking a tabu have been defined [16], and they are commonly known as aspiration criteria. The simplest and most commonly used aspiration criterion consists of allowing a move (even if it is tabu) if it results in a solution with an objective value better than that of the current best-known solution (since the new solution has obviously not been previously visited).

The basic TS algorithm, considering the minimization of the objective function f, works as follows:

1. randomly select an initial solution i in the search space S, and set $i^* = i$ and $k = 0$, where i^* is the best solution so far, and k the iteration counter;
2. set $k = k+1$ and generate the subset V of the *admissible* neighborhood solutions of i (i.e., non-*tabu* or allowed by aspiration);
3. choose the best j in V and set $i = j$;
4. if $f(i) < f(i^*)$, then set $i^* = i$;
5. update the *tabu* and the aspiration conditions;
6. if a stopping condition is met then stop; else go to Step 2.

The most commonly adopted conditions to end the algorithm are when the number of iterations (K) is larger than a maximum number of allowed iterations, or if no changes to the best solution have been performed in the last N iterations (as it is in our tests).

Specifically, in our experiments, we modeled for both HC and TS the possible valid solutions in the search space as a binary adjacency matrix describing acyclic directed graphs. The starting point of the search is the empty graph (i.e. without any edge), and the search is stopped when the current best solution cannot be improved with any move in its neighborhood.

The algorithms then consider a set of possible solutions (i.e., all the directed acyclic graphs) and navigate among them by means of 2 moves: insertion of a new edge or removal of an edge currently in the topology. We also recall that in the literature many alternatives are proposed to navigate the search space when learning the structure of Bayesian Networks, see [17, 18]. But, for the purpose of this work, we preferred to stick with the classical (and simpler) ones.

3.3. Genetic Algorithms

Genetic Algorithms (GAs) [19] are a class of computational models that mimic the process of natural evolution. GAs are often viewed as function optimizers although the range of problems to which GAs have been applied is quite broad. Their peculiarity is that potential solutions that undergo evolution are represented as fixed length strings of characters or numbers. Generation by generation, GAs stochastically transform sets (also called *populations*) of candidate solutions into new, hopefully improved, populations of solutions, with the goal of finding one solution that suitably solves the problem at hand. The quality of each candidate solution is expressed by using a user-defined objective function called *fitness*. The search process of GAs is shown in Figure 1.

To transform a population of candidate solutions into a new one, GAs make use of particular operators that transform the candidate solutions, called genetic operators: *crossover* and *mutation*. Crossover is traditionally used to combine the genetic material of two solutions (parents) by swapping a part of one individual with a part of the other. On the other hand, mutations introduce random changes in the structures of the solutions. In order to be able to use GAs to solve a given optimization problem, candidate solutions must be encoded into individuals (solutions) and, often, also the genetic operators (crossover and mutation) must be adapted.

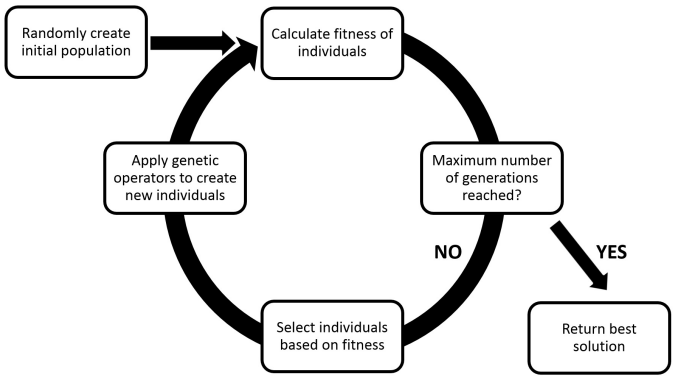

Figure 1: Typical workflow of a GA algorithm.

GAs have been widely used to learn the BN structure considering the search space of the DAGs. In the large majority of the works [20, 21, 22, 23, 24, 25], the GA encodes the connectivity matrix of the BN structure in its individuals. This is the same approach used in our study. Our GA follows the common structure:

1. Generate the initial population
2. Repeat the following operations until the total number of generations is reached:
   (a) Select a population of parents by repeatedly applying the parent selection method
   (b) Create a population of offspring by applying the crossover operator to each set of parents
   (c) For each offspring, apply one of the mutation operators with a given mutation probability
   (d) Select the new population by applying the survivor selection method
3. Return the best performing individual

Unless stated otherwise, the following description is valid for both GA variants (discrete and continuous). The initialization method and the variation operators used in our GA ensure that every individual is valid, i.e., each individual is guaranteed to be an acyclic graph. The initialization method creates individuals with exactly $N/2$ connections between variables, where N is the total number of variables. A connection is only added if the graph

remains acyclic. The nodes being connected are selected with uniform probability from the set of variables. In the continuous variant, the value associated with each connection is randomly generated between 0.0 and 1.0, with uniform probability. The crossover operates over two parents and returns one offspring. The offspring is initially a copy of the first parent. Afterward, a node with at least one connection is selected from the second parent and all the valid connections starting from this node are added to the offspring. Three mutation operators are considered in our GA implementation: add connection mutation, remove connection mutation, and perturbation mutation. The first two are applied in both GA variants, while the perturbation mutation is only applied in the continuous variant. The offspring resulting from the add connection mutation operator differs from the parent in having an additional connection. The newly connected nodes are selected with uniform probability. In the continuous variant, the value associated with the new connection is randomly generated between 0.0 and 1.0 with uniform probability. Similarly, the offspring resulting from the remove connection mutation operator differs from the parent in having one less connection. The nodes to be disconnected are selected with uniform probability. The perturbation mutation operator applies Gaussian perturbations to the values associated with at most $N/2$ connections. The total number of perturbations may be less than $N/2$ if the individual being mutated has less than $N/2$ connections. Each value is perturbed following a Gaussian distribution having mean 0.0 and standard deviation 1.0. Any resulting value below 0.0 or above 1.0 is bounded to 0.0 or 1.0, respectively. Regardless of the GA variant, mutation is applied with a probability of 0.25. When a mutation is being applied, the specific mutation operator is selected with uniform probability from the available options (two in the discrete variant, and three in the continuous variant).

In terms of parameters, a population of size 100 is used, with the evolutionary process being able to perform 100 generations. Parents are selected by applying a tournament of size 3. The survivor selection is elitist in the sense that it ensures that the best individual survives to the next generation.

4. Results and Discussion

We made use of a large set of simulations on randomly generated datasets with the aim of assessing the characteristics of the state-of-the-art techniques for structure learning of BNs.

We generated data for both the case of discrete (2 values, i.e., 0 or 1) and continuous (we used 4 levels, i.e., 1, 2, 3, and 4) random variables. For each of them, we randomly generated both the structure (i.e, 100 weakly connected DAGs having 10 and 15 nodes) and the related parameters (we generated random values in $(0, 1)$ for each entry of the conditional probability tables attached to the network structure) to build the simulated BNs. We also considered 3 levels of density of the networks, namely 0.4, 0.6, and 0.8 of the complete graph. For each of these scenarios, we randomly sampled from the BNs different datasets of sample sizes, parametrized based on the number of nodes. Specifically, for networks of 10 variables, we generated datasets of 10, 50, 100, and 500 samples, while for 15 variables, we considered datasets of 15, 75, 150, and 750 samples. Furthermore, we also considered noise in the samples as a set of random entries (both false positives and false negatives) in the dataset. To this extent, we considered noise free dataset (noise rate equals 0%) and dataset with an error rate of 10% and 20%. In total this led us to a total number of 14400 random datasets.

For all of them, we considered both the constraint-based and the score-based approaches to structural learning. From the former category of methods, we considered the PC algorithm [7] and the IAMB [10]. We recall that these methods return a partially directed graph, leaving undirected the arcs that are not unequivocally directable. In order to have a fair comparison with the score-based method which returns DAGs, we randomly resolved the ambiguities, by generating random solutions (i.e., DAGs) consistent with the statistical constraints by PC and IAMB (that is, we select a random direction for the undirected arcs).

Moreover, among the score-based approaches, we consider 3 maximum likelihood scores, namely log-likelihood [11], BIC [12], and AIC [26] (for continuous variables we used the corresponding Gaussian likelihood scores). For all of them, we repeated the inference on each configuration by using HC, TS, and GA as search strategies.

This led us to 11 different methods to be compared, for a total of 158400 independent experiments. To evaluate the obtained results, we considered both false positives, FP (i.e., arcs that are not in the generative BN but that are inferred by the algorithm) and false negatives, FN (i.e., arcs that are in the generative BN but that are not inferred by the algorithm). Also, with TP (true positives) being the arcs in the generative model and TN (true negatives) the arcs that are not in the generative model, for all the

experiments we considered the following metrics:

$$Precision = \frac{TP}{TP + FP} \qquad Recall = \frac{TP}{TP + FN}$$

$$Specificity = \frac{TN}{TN + FP} \qquad Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

4.1. Results of the Simulations

In this Section, we comment on the results of our simulations. As anticipated, we computed precision, recall (sensitivity), and specificity, as well as accuracy and *Hamming* distance, to assess the performances and the underfit/overfit trade-off of the different approaches. Overall, from the obtained results, it is straightforward to notice that methods including more edges in the inferred networks are also more subject to errors in terms of accuracy, which may also resemble a bias of this metric that tends to penalize solutions with false positive edges rather than false negatives. On the other hand, since a typical goal of the problems involving the inference of BNs is the identification of novel relations (that is, proposing novel edges in the network), underfitting approaches could be more effective in terms of accuracy, but less useful in practice.

With a more careful look, the first evidence we obtained from the simulations is that the two parameters with the highest impact on the inferred networks are the density and the number of nodes (i.e., network size). For this reason, we first focused our attention on these two parameters, and we analyzed how the different tested combinations of methods and scores behave.

As shown in Figure 2, all the approaches seem to perform better when dealing with low-density networks, in fact, for almost all the methods the accuracy is higher for density equal to 0.4, while it is lower for density equal to 0.8. Since each edge of the network is a parameter that has to be "learned", it is reasonable to think that, the more edges are present in the BN, the harder becomes the problem to be solved when learning it. Moreover, we can also observe that the results in terms of accuracy obtained for BN of discrete variables are slightly better than those achieved on the continuous ones. To this extent, the only outlier is the loglik score combined with Hill Climbing and with Tabu Search on datasets with continuous variables, for which the trend is the opposite, w.r.t. that of all the other approaches. In fact, in both these cases, the accuracy is higher on high-density datasets, and this is likely

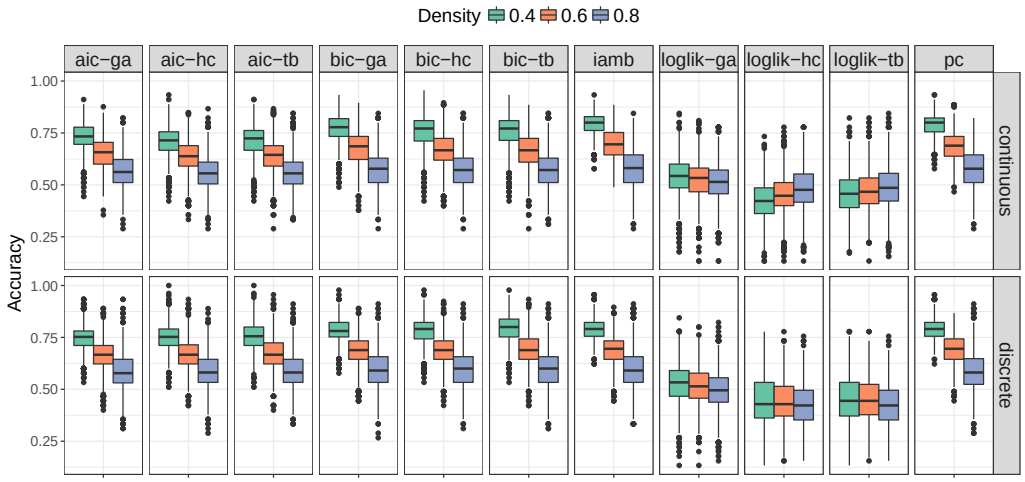

Figure 2: Boxplots showing the accuracy of the results obtained by the 11 tests approaches on the discrete and continuous datasets.

due to a very high overfit of the approach. It is interesting to notice that Genetic Algorithms (combined with the loglik score) is less affected by this problem, and this trend will also be shown in the next analyses.

In addition to the accuracy, we also computed the *Hamming* distance between the reconstructed solutions and the BNs used to generate the data, in order to quantify the errors of the inference process. This analysis showed that, besides the network density, also the number of nodes (parameters to be learned) influences the results, as reported in Figure 3. It is interesting to observe that, in the adopted experimental setup, we set the number of samples to be proportional to the number of nodes of the network, that is, for 10 nodes, in the same configurations we have a lower number of samples, compared to the ones for 15 nodes. From a statistical point of view alone, we would expect the problem to be easier when having more samples to build the BN, since this would lead to more statistical power, and, intuitively, should compensate for the fact that with 15 nodes we have more parameters to learn than the ones for the case of 10 nodes. While this may be the case (in fact we have similar accuracy for both 10 and 15 nodes), we also observe a constantly higher *Hamming* distance with more nodes. In fact, when dealing with more variables we observed a shift in the performance, that is, even when the density is low, we observe more errors manifesting with higher

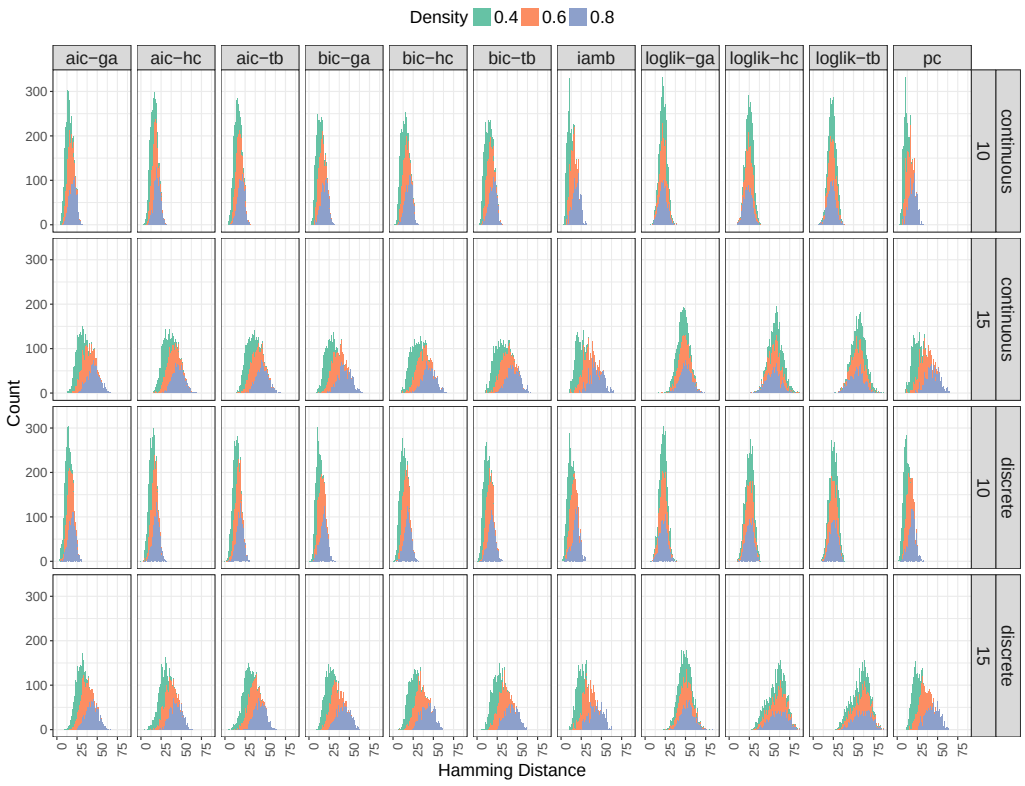

Figure 3: Bar plots showing the *Hamming* distances of the solutions obtained by the 11 tested approaches, grouped by type of dataset (discrete and continuous) and number of nodes of the network (10 and 15), while colors represent densities (0.4, 0.6, and 0.8).

values in terms of *Hamming* distance. This is due to the fact that, when increasing the number of variables, we also increase the complexity of the solutions.

A further analysis we performed was devoted to assessing the impact of both overfit and underfit. As it is possible to observe from the plots in Figure 4 and Figure 5, we obtain two opposite behaviors, namely, the combinations of scores and search strategies show evident different trends in terms of *sensitivity* (*recall*), *specificity*, and *precision*. In details, *iamb* and *pc* tend to underfit, since they both produce networks with consistently low density. While they achieve similar (high) results in terms of accuracy (see Figure 2), their trend toward underfit does not make them suited for the identification of novel edges, rather being more indicated for descriptive

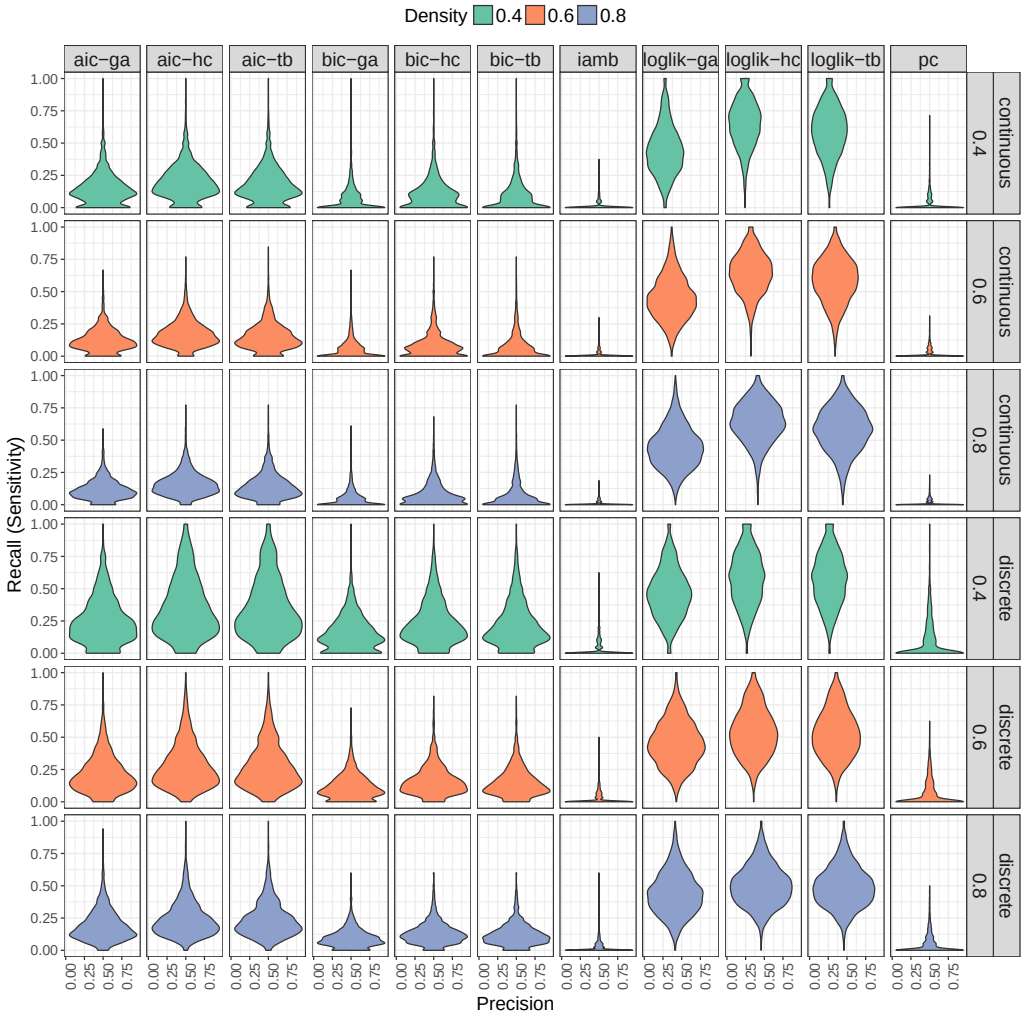

Figure 4: Violin plots showing Precision and Recall of the solutions obtained by the 11 tested approaches, grouped by the type of dataset (discrete and continuous) and density values (0.4, 0.6, and 0.8).

purposes. On the other hand, the *loglik* score, (mostly) independently from the adopted search technique, consistently overfits. Both these behaviors can be observed in the violin plots of Figure 4 and Figure 5: for each density, we can observe that iamb/pc results have very low *recall* values, while the distribution of the results of the loglik score is centered on higher values.

The other two scores, i.e., AIC and BIC, present a better trade-off between

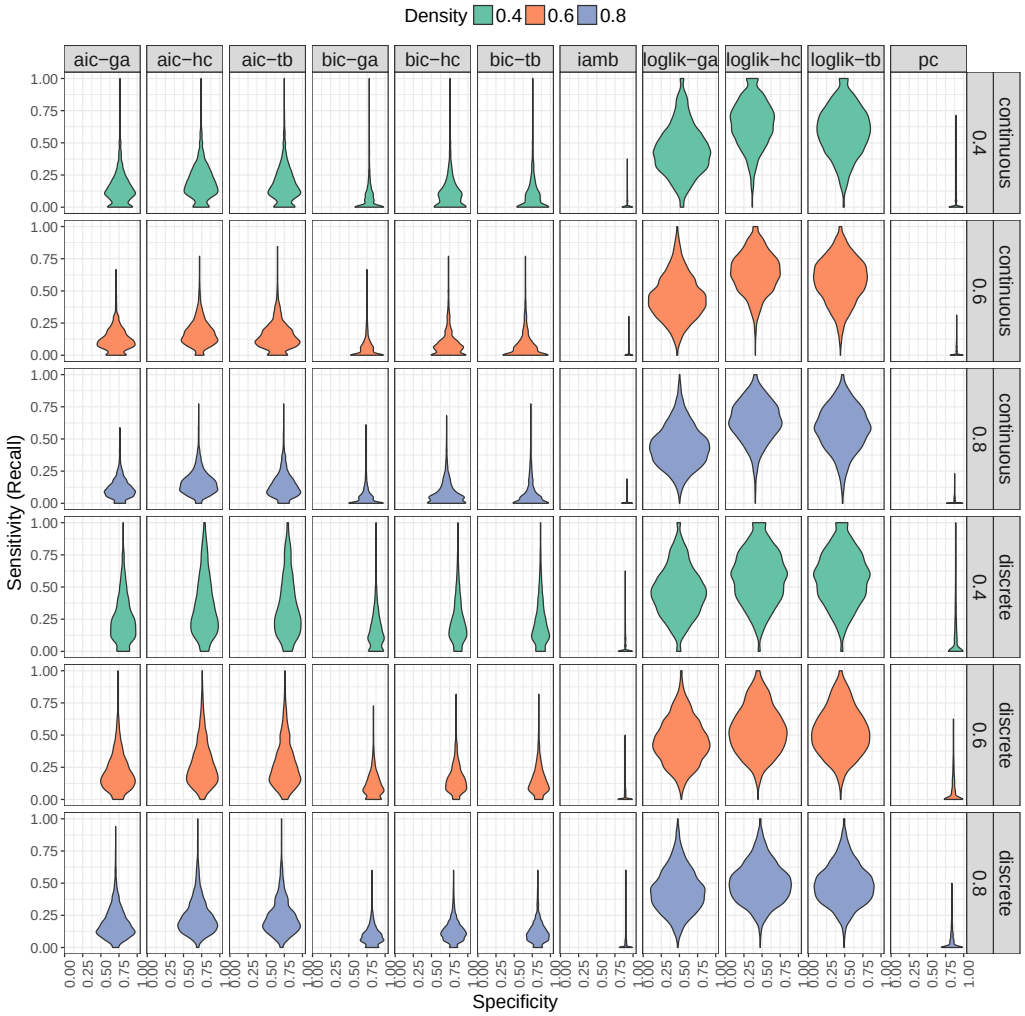

Figure 5: Violin plots showing Specificity and Sensitivity of the solutions obtained by the 11 tested approaches, grouped by the type of dataset (discrete and continuous) and density values (0.4, 0.6, and 0.8).

overfit and underfit, with BIC being less affected by the overfit, while AIC reconstructing slightly denser networks. Once again, this trade-off between the two regularizators is well-known in the literature and points to AIC being more suited for predictions while BIC for descriptive purposes.

Another relevant result of our analysis is the characterization of the performance of Genetic Algorithms in terms of *sensitivity* (*recall*). Specifically,

it must be noticed that for all the three regularizators, Genetic Algorithms achieve similar results of both *precision* and *specificity* if compared with both Hill Climbing and Tabu Search, but in terms of *sensitivity* (*recall*) it presents reduced overfit. In fact, as it is possible to observe in the plots of Figure 4 and Figure 5 for each score (i.e., loglik, AIC, and BIC), the *sensitivity* results of Genetic Algorithms are lower than that of both Hill Climbing and Tabu Search, highlighting the reduced impact of overfit.

In summary, we can draw the following conclusions.

- (*i*) dense networks are harder to learn, see Figure 2;
- (*ii*) the number of nodes effects the complexity of the solutions, leading to higher *Hamming* distance (number of errors) even if with similar accuracy, see Figure 2 and Figure 3;
- (*iii*) networks with continuous variables are harder to learn compared to the binary ones, see Figure 2;
- (*iv*) *iamb* and *pc* algorithms tend to underfit, while *loglik* overfits, see Figure 4 and Figure 5;
- (*v*) Genetic Algorithms tend to reduce underfit, see Figure 4 and Figure 5.

We conclude the section by providing $p-$values by MannWhitney U test [27] in support of such claims in Table 4.1. The claim for greater and less values, i.e., accuracy, are performed with the one-tail test alternatives. The tests are performed on all the configurations of our simulations.

5. Conclusions

Bayesian Networks are a widespread technique to model dependencies among random variables. A challenging problem when dealing with BNs is the one of learning their structures, i.e., the statistical dependencies in the data, which may sometime pose a serious limit to the reliability of the results.

Despite their extensive use in a vast set of fields, to the best of our knowledge, a quantitative assessment of the performance of different state-of-the-art methods to learn the structure of BNs has never been performed. In this work, we aim at going in the direction of filling this gap and we presented a study of different state-of-the-art approaches for structural learning

Comparison	**Metric**	**Test**	$p-$**value**	**Mean** 1	**Mean** 2
density 0.4/0.8	accuracy	greater	$< 2.2e-16$	0.68	0.55
nodes 10/15	accuracy	less	$< 2.2e-16$	0.61	0.62
nodes 10/15	Hamming	less	$< 2.2e-16$	17.40	39.54
discr./contin.	accuracy	greater	$< 2.2e-16$	0.62	0.61
iamb/hc loglik	sensitivity	less	$< 2.2e-16$	0.02	0.59
iamb/hc loglik	specificity	greater	$< 2.2e-16$	0.98	0.38
ga/hc	sensitivity	less	$< 2.2e-16$	0.24	0.32
ga/hc	specificity	greater	$< 2.2e-16$	0.79	0.71

Table 1: Summary of the major findings. The results of the MannWhitney U test [27] (one-tail test alternatives) in support of the results of our simulations are shown. The tests are performed on all the settings, with the comparisons as described in the first column of the table. We also show the considered metric, the adopted alternative of the one-tail test, the obtained $p-$value, and the mean of the two compared distributions of results.

of Bayesian Networks on simulated data, considering both discrete and continuous variables.

To this extent, we investigated the characteristics of 3 different likelihood score, combined with 3 commonly used search strategies schemes (namely, Genetic Algorithms, Hill Climbing, and Tabu Search), as well as 2 constraint-based techniques for the BN inference (i.e., iamb and pc algorithms).

Our analysis showed the factors leading to higher effects on the performance, i.e., density, number of variables, and variable type (discrete vs continuous). As expected, these settings effect the number of parameters to be learned, hence complicating the optimization task. Furthermore, we also discussed the overit/underfit trade-off of the different tested techniques with the constraint-based approaches showing trends toward underfitting, and the loglik score showing high overfit. Interestingly, in all the configurations, Genetic Algorithms showed evidence of reducing the overfit, leading to denser structures.

Overall, we place our work as a starting effort to better characterize the task of learning the structure of Bayesian Networks from data, which may lead in the future to a more effective application of this approach.

References

References

[1] S.-S. Leu, Q.-N. Bui, Leak prediction model for water distribution networks created using a bayesian network learning approach, Water Resources Management 30 (2016) 2719–2733.

[2] F. Dondelinger, S. Lèbre, D. Husmeier, Non-homogeneous dynamic bayesian networks with bayesian regularization for inferring gene regulatory networks with gradually time-varying structure, Machine Learning 90 (2013) 191–230.

[3] W. C. Young, A. E. Raftery, K. Y. Yeung, Fast bayesian inference for gene regulatory networks using scanbma, BMC Systems Biology 8 (2014) 47.

[4] W. Buntine, A guide to the literature on learning probabilistic networks from data, IEEE Transactions on knowledge and data engineering 8 (1996) 195–210.

[5] R. Daly, Q. Shen, S. Aitken, Learning bayesian networks: approaches and issues, The Knowledge Engineering Review 26 (2011) 99–157.

[6] P. Larrañaga, H. Karshenas, C. Bielza, R. Santana, A review on evolutionary algorithms in bayesian network learning and inference tasks, Information Sciences 233 (2013) 109 – 125.

[7] P. Spirtes, C. N. Glymour, R. Scheines, Causation, prediction, and search, MIT press, 2000.

[8] D. M. Chickering, Learning bayesian networks is np-complete, in: Learning from data, Springer, 1996, pp. 121–130.

[9] D. M. Chickering, D. Heckerman, C. Meek, Large-sample learning of bayesian networks is np-hard, Journal of Machine Learning Research 5 (2004) 1287–1330.

[10] I. Tsamardinos, C. F. Aliferis, A. R. Statnikov, E. Statnikov, Algorithms for large scale markov blanket discovery., in: FLAIRS conference, volume 2, pp. 376–380.

[11] D. Koller, N. Friedman, Probabilistic graphical models: principles and techniques, MIT press, 2009.

[12] G. Schwarz, et al., Estimating the dimension of a model, The annals of statistics 6 (1978) 461–464.

[13] C.-R. Hwang, Simulated annealing: theory and applications, Acta Applicandae Mathematicae 12 (1988) 108–111.

[14] F. Glover, Tabu search-part i, ORSA Journal on computing 1 (1989) 190–206.

[15] M. Gendreau, J.-Y. Potvin, Handbook of metaheuristics, volume 2, Springer, 2010.

[16] F. Glover, Tabu search: A tutorial, Interfaces 20 (1990) 74–94.

[17] D. M. Chickering, Learning equivalence classes of bayesian-network structures, Journal of machine learning research 2 (2002) 445–498.

[18] M. Teyssier, D. Koller, Ordering-based search: A simple and effective algorithm for learning bayesian networks, arXiv preprint arXiv:1207.1429 (2012).

[19] D. E. Goldberg, J. H. Holland, Genetic algorithms and machine learning, Machine learning 3 (1988) 95–99.

[20] C. Cotta, J. Muruzábal, Towards a more efficient evolutionary induction of bayesian networks, in: International Conference on Parallel Problem Solving from Nature, Springer, pp. 730–739.

[21] S. Van Dijk, D. Thierens, L. C. Van Der Gaag, Building a ga from design principles for learning bayesian networks, in: Genetic and Evolutionary Computation Conference, Springer, pp. 886–897.

[22] R. Etxeberria, P. Larrañaga, J. M. Picaza, Analysis of the behaviour of genetic algorithms when learning bayesian network structure from data, Pattern Recognition Letters 18 (1997) 1269–1273.

[23] K.-J. Kim, J.-O. Yoo, S.-B. Cho, Robust inference of bayesian networks using speciated evolution and ensemble, in: International Symposium on Methodologies for Intelligent Systems, Springer, pp. 92–101.

[24] M. Mascherini, F. M. Stefanini, M-ga: A genetic algorithm to search for the best conditional gaussian bayesian network, in: International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC'06), volume 2, IEEE, pp. 61–67.

[25] P. Larrañaga, M. Poza, Y. Yurramendi, R. H. Murga, C. M. H. Kuijpers, Structure learning of bayesian networks by genetic algorithms: A performance analysis of control parameters, IEEE transactions on pattern analysis and machine intelligence 18 (1996) 912–926.

[26] H. Akaike, Information theory and an extension of the maximum likelihood principle, in: Selected Papers of Hirotugu Akaike, Springer, 1998, pp. 199–213.

[27] H. B. Mann, D. R. Whitney, On a test of whether one of two random variables is stochastically larger than the other, The annals of mathematical statistics (1947) 50–60.